

A Self-Consistent Dealer-Flow Field: Nonlinear-Schrödinger Modelling of Spot-Options Feedback

Maksim Aleksandrovich Andreev, Joseph Howden
Quark Quantitative Research

Abstract

The Black–Scholes (1973) framework is a one-way map: it prices an option as a function of an exogenously given underlying, under the assumption that hedging the option does not move the underlying. That assumption is structurally violated in the regime that now dominates index microstructure, where dealers hold large, concentrated options inventories and their delta-hedging of that inventory exerts a measurable mechanical force on the spot (Gârleanu, Pedersen & Poteshman, 2009; Ni, Pearson & Poteshman, 2005; Barbon & Buraschi, 2020). We formalise the omitted back-reaction. Following the realist (Bohmian) reading in which the price has a definite value carried by a real present field rather than a probability cloud (Bohm, 1952), we model the underlying as a particle guided by a self-consistent field that its own options positioning creates. Dealer delta-hedging exerts a measurable price-impact drift on the underlying, sourced by the dealer-gamma potential $V(S) = -\int \Gamma_S dS$ (Kyle, 1985; Avellaneda & Lipkin, 2003); by the Madelung (1927)–Nelson (1966) transform — an exact change of variables, not an analogy — this diffusion is mathematically identical to a Schrödinger equation carrying V as its external potential, with Black–Scholes the zero-potential (free-particle) limit. Because the dealer position depends on where the price sits, V is density-dependent and the diffusion is of McKean–Vlasov (mean-field) type, whose Madelung representation is a *nonlinear* Schrödinger equation; for a local density-dependence the nonlinearity is the Gross–Pitaevskii term $g|\psi|^2\psi$. The transforms are theorems and V is measurable from the book; the single modelling choice — the functional form of the density-dependence (the self-focusing factor) — is precisely what the forward test validates. The scalar Gamma Exposure statistic (GEX) used by practitioners is recovered as the *linear* layer; the novel object is the *nonlinear self-focusing* layer, in which the pinning force intensifies as the spot localises near a high-gamma strike and as charm sharpens the landscape into expiry. We give the field reconstruction from a listed options chain and a pre-registered forward-validation protocol (Harvey, Liu & Zhu, 2016) whose decisive test is whether the self-focusing increment forecasts short-horizon pinning *beyond* linear GEX. This is a methods and protocol paper; empirical results are reported separately and only on forward, point-in-time data, per the validation posture of Andreev & Howden (2026e).

1. Introduction

The Black–Scholes–Merton model treats the underlying price process as exogenous: the option is a derivative claim whose value is fixed by no-arbitrage replication against an underlying that the replication itself is assumed not to perturb. This is an excellent approximation when the hedging demand is small relative to the liquidity of the underlying. It fails when it is not. The documented exceptions are no longer marginal: dealer inventories in index and single-name options are large and concentrated, and the delta-hedging of that inventory is a first-order driver of realised spot dynamics — clustering of expirations at high open-interest strikes (Ni, Pearson & Poteshman,

2005), demand-pressure effects on the option surface itself (Gârleanu, Pedersen & Poteshman, 2009), and the amplification or suppression of realised volatility as a function of the aggregate dealer gamma position (Barbon & Buraschi, 2020). Avellaneda & Lipkin (2003) gave the canonical model of the simplest such effect — expiration pinning — as a market-induced mechanism arising directly from hedging feedback.

What is missing is a unified field description of the two-way coupling, of which pinning is one limit. The practitioner statistic in common use, dealer Gamma Exposure (GEX), is a single scalar: the net dollar gamma the dealer must re-hedge per unit move in the spot, evaluated at the current spot. It captures the *sign* of the feedback (long-gamma dealers dampen moves; short-gamma dealers amplify them) but treats the

landscape as static and the response as linear. It does not encode how the force changes as the price moves through the landscape, nor how the landscape itself sharpens as time-decay concentrates gamma into the near-dated strikes.

Contribution. We give a self-consistent field model of the spot–options feedback. (i) We derive the price’s effective dynamics under dealer hedging — a measurable price-impact drift sourced by the dealer-gamma potential V — and show, via the Madelung–Nelson transform, that the diffusion is mathematically identical to a Schrödinger equation carrying V (scalar GEX is the linear layer; Black–Scholes the $V = 0$ limit). (ii) We show the nonlinearity is *forced*, not added: because the dealer position — and hence V — depends on the price density, the diffusion is mean-field (McKean–Vlasov) and its Madelung representation carries a Gross–Pitaevskii self-interaction (a $|\psi|^2\psi$ term), a change of the governing equation rather than a parameterised extension of Black–Scholes. (iii) We give an explicit reconstruction of the field from a listed chain and a pre-registered forward test whose decisive bar is incremental predictive power of the self-focusing term over linear GEX — so that a confirmation is evidence for the nonlinear structure specifically, not a re-derivation of dealer gamma. The object is a microstructure / execution model; it is never an input to portfolio allocation.

2. The One-Way Assumption and Its Failure

In the Black–Scholes derivation the underlying follows an exogenous diffusion $dS = \mu S dt + \sigma S dW$, and the option value $C(S,t)$ solves the backward equation obtained by forming a locally riskless replicating portfolio. The replication is one-way by construction: the hedger’s trades in S are assumed to leave the law of S unchanged. Under the standard log-transform the pricing equation is the heat equation, i.e. an imaginary-time free-particle Schrödinger equation with vanishing potential; the implied-volatility smile is the leading empirical signal that the true potential is not zero (Baaquie, 2004).

The failure mode of interest here is not the smile (a static mis-specification of the marginal) but the *dynamic* back-reaction. Let the dealer aggregate carry signed net dollar gamma $\Gamma_\$(S)$ (the convention made precise in §3.2). To remain delta-neutral the dealer trades $-\Gamma_\$(S) \cdot dS$ of the underlying for a move dS . When the dealer is long gamma ($\Gamma_\$ > 0$) this is a contrarian flow (sell into rallies, buy into dips) that damps realised volatility and pins the spot toward the gamma-weighted centre; when short gamma it is a momentum flow that amplifies moves and can drive squeezes. This is a force on S that depends on S — precisely the term Black–Scholes assumes away; §3 makes it quantitative.




Figure 1. Schematic of the two hedging regimes. The dealer book's re-hedging flow per price move dS is $-\Gamma_S \cdot dS$. (a) When dealers are net long gamma ($\Gamma_S > 0$) the flow opposes the move — rallies are met by selling and dips by buying — damping realised volatility and drawing the spot toward the gamma-weighted centre (the pinning regime; the pin K^* and the potential well are made precise in §3.2). (b) When dealers are net short gamma the same delta-neutrality rule chases the move — rallies are bought and dips sold — amplifying realised volatility and feeding squeezes. Axes are unscaled: the diagram is a schematic of the mechanism, not data.

3. The Self-Consistent Field

3.1 Derivation: from hedging feedback to a (nonlinear) Schrödinger equation

The correspondences below are exact transforms and a measurable mechanism, not analogies. We build the field in four steps and flag at the end the single component that is a modelling choice rather than a derived identity.

(1) The hedging force is measurable. A dealer aggregate carrying net dollar gamma $\Gamma_S(S,t)$ must, to stay delta-neutral, trade $-\Gamma_S(S,t) \cdot dS$ of the underlying per move dS . Through a linear price-impact coefficient λ (Kyle, 1985; Almgren & Chriss, 2001) this re-hedging flow is an effective drift on the underlying, $b(S,t) = -\lambda \cdot \partial V / \partial S$, with dealer-gamma potential $V(S,t) = -\int \Gamma_S(S',t) dS'$. Avellaneda & Lipkin (2003) derive exactly this drift for the pinning case. Both b and V are reconstructed from the observable options book — this is the empirical content of the paper.

(2) The diffusion is exactly a Schrödinger equation.

The price obeys $dS = b(S,t) dt + \sigma dW$. By the Madelung (1927)–Nelson (1966) transform — the exact change of variables $\psi = \sqrt{\rho} \cdot \exp(i\phi/\sigma^2)$, not a resemblance — this is mathematically identical to $i\sigma^2 \partial_t \psi = [-(\sigma^4/2) \partial_S^2 + V(S,t)] \psi$. V is the external potential in the Hamiltonian by construction of the transform. The Madelung quantum potential $Q = -(\sigma^4/2)(\partial_S^2 \sqrt{\rho})/\sqrt{\rho}$ is a *distinct*, density-intrinsic term and is not the dealer-gamma landscape — an earlier informal identification of the two is corrected here.

(3) Black–Scholes is the $V = 0$ limit. With no dealer back-reaction ($\lambda = 0$ or $\Gamma_S \equiv 0$) the potential vanishes and the equation reduces to the free-particle (heat-equation) form Black–Scholes inherits — recovered exactly as a special case, not approximated away.

(4) The nonlinearity is forced, not chosen. The dealer aggregate position, and hence V , depends on where price and customer demand actually sit, so $V = V[\rho]$ is a functional of the price density. The diffusion is then of McKean–Vlasov (mean-field) type and its Madelung representation is a *nonlinear* Schrödinger equation; for a local density-dependence $V[\rho] \supset g \cdot \rho$

$= g|\psi|^2$ the nonlinearity is exactly the Gross–Pitaevskii term $g|\psi|^2\psi$ (Gross, 1961; Pitaevskii, 1961). The nonlinear-Schrödinger structure is therefore a theorem about density-dependent drift, not a stylistic choice.

Derived vs modelled. Steps (2)–(4) are exact (the Madelung–Nelson transform and the McKean–Vlasov $\rightarrow$ nonlinear-Schrödinger correspondence are theorems); in step (1), b and V are empirically measurable from the book. The one component that is a *modelling choice*, not a derived identity, is the specific functional form of the density-dependence $V[\rho]$ — equivalently the self-focusing factor Φ of §3.4. That form is the hypothesis the forward test of §5 exists to validate or reject; it is not asserted as certain.




Figure 2. Schematic of the self-consistent feedback loop derived in §3.1. A spot move dS forces the delta-neutral dealer book to trade $-\Gamma_S \cdot dS$ of the underlying; through the linear price-impact coefficient λ (Kyle, 1985) that aggregate flow is a drift $b = -\lambda \cdot \partial V / \partial S$ on the spot, which moves the price again — closing the loop the Black–Scholes replication argument assumes open. Because the dealer position, and hence the potential V , depends on where the price density ρ actually sits (centre box), the loop is a mean-field (McKean–Vlasov) system whose Madelung representation is a nonlinear Schrödinger equation carrying the Gross–Pitaevskii self-interaction $g|\psi|^2\psi$: as the spot localises on a concentrated-gamma strike the pull deepens — the self-focusing channel parameterised in §3.4. Schematic only; no data.

3.2 The dealer-gamma potential (the external potential V)

Sections 3.2–3.4 reconstruct the three ingredients of the §3.1 derivation — the potential V , its linear gradient, and the self-focusing factor — explicitly from a listed options chain. For a chain with strikes K ,

per-strike open interest $OI(K)$ and per-strike Black–Scholes gamma $\gamma(S, K, \tau)$, the dealer dollar-gamma density (in the standard convention that dealers are long call gamma and short put gamma relative to customer demand) is

$$\Gamma_{\$}(S) = \sum_K s(K) \cdot OI(K) \cdot \gamma(S, K, \tau) \cdot 100 \cdot S^2 \cdot 0.01$$

with $s(K)$ the dealer sign at strike K . The scalar GEX is $\Gamma_{\$}$ evaluated at the current spot. This external potential is the V derived in §3.1, the integral of the

restoring force, $V(S) \propto -\int \Gamma_{\$}(S') dS'$; in the long-gamma regime V has a well whose minimum is the pin K^* , taken as the $|\Gamma_{\$}|$ -weighted centroid of the near-the-money strikes, and the gamma-flip level S_{flip} is the spot at which $\Gamma_{\$}$ changes sign (the boundary between the stabilising and destabilising regimes).




Figure 3. The objects of §3.2 on a real chain: net dealer dollar-gamma profile by strike for SPY, reconstructed from the 18 July 2026 record of the daily gamma-density snapshot log (per-strike open interest and implied volatilities across listed expiries to seventy days, each strike's dollar gamma signed by the dealer convention above and smeared with a kernel of one percent of spot onto a fixed $\pm 15\%$ moneyness grid). Marked verticals: the spot (743.3), the gamma flip (756.1, the strike at which the net profile changes sign), and the pin K^* (733.5, the absolute-gamma-weighted centroid of the $\pm 15\%$ window). On this snapshot the book sits in the short-gamma regime at the spot — net dealer gamma there is $-\$5.1\text{B}$ per 1% move, dominated by the put wall just below the money — with the stabilising long-gamma lobe above the flip. Descriptive exhibit only: no outcome statistic is computed.

3.3 The guidance drift (linear layer)

The hedging drift b derived in §3.1 is, to linear order in the displacement from the pin, the deterministic pull the spot inherits from the dealer potential; it points toward (long gamma) or away from (short gamma) the pin with a magnitude set by the scalar gamma:

$$v_{\text{lin}}(S) = \kappa \cdot [\Gamma_{\$}(S) / \Gamma_{\text{scale}}] \cdot (K^* - S)/S$$

where Γ_{scale} normalises by the in-window dealer dollar-gamma so that v is a dimensionless fractional drift, and κ is absorbed into the regression slope of the empirical test (§5). This linear layer is exactly the

information content of practitioner GEX: a signed, distance-weighted pull whose strength is the scalar gamma at the current spot.

3.4 The self-interaction (nonlinear layer)

The linear layer treats the landscape as fixed and the response as proportional to the scalar $\Gamma_{\$}(S)$. But the gamma density is sharply peaked at-the-money, so the force the spot actually feels depends on *where the spot sits* within the landscape; and charm ($\partial\Gamma/\partial t$) concentrates that density into the near-dated strikes as $\tau \rightarrow 0$. Both effects make the effective potential depend on the present configuration of the field — the density-dependence $V[p]$ that, per the McKean–

Vlasov $\rightarrow$ nonlinear-Schrödinger correspondence of §3.1(4), enters as the Gross–Pitaevskii $|\psi|^2\psi$ self-interaction (Gross, 1961; Pitaevskii, 1961). We parameterise the self-focused drift as the linear drift modulated by a self-interaction factor Φ — this functional form is the one modelling choice flagged in §3.1, validated by the forward test, not a derived identity:

$$v_{nl}(S, \tau) = v_{lin}(S) \cdot \Phi_{local}(S) \cdot \Phi_{expiry}(\tau)$$

$$\Phi_{local}(S) = |\gamma| \text{ at the strike nearest } S, \text{ relative to the window-mean } |\gamma|$$

$$\Phi_{expiry}(\tau) = (\tau_{ref} / \tau)^p, \quad 0 < p \leq 1$$

Φ_{local} deepens the local well when the spot sits on a concentrated-gamma strike (self-trapping); Φ_{expiry} sharpens the pin into expiry (the charm channel). Both are bounded by frozen clamps to prevent a single thin-OI strike or the $\tau \rightarrow 0$ limit from diverging. This is self-focusing in the nonlinear-Schrödinger system of §3.1: the price digs its own potential well and is trapped by it, the same mechanism that produces solitons and self-trapped states in Gross–Pitaevskii systems.

Linear vs nonlinear, made testable. The linear layer uses only the scalar $\Gamma_{\S}(S_{now})$; the nonlinear layer multiplies it by $\Phi(S, \tau)$. The *increment* $v_{nl} - v_{lin}$ is the entire empirical content of the nonlinear structure. If that increment carries no predictive power beyond v_{lin} , the nonlinear layer is decoration and the model reduces to GEX; if it does, the self-consistent field is doing work the scalar statistic cannot.

4. Relation to Black–Scholes and to the Stochastic-Mechanical Architecture

By the Madelung–Nelson transform of §3.1, Black–Scholes is exactly the linear Schrödinger equation with vanishing external potential: the wavefunction evolves in a fixed background and does not act back on it. The dealer-flow field is the nonlinear Schrödinger / Gross–Pitaevskii equation the same transform yields once the potential is non-zero and density-dependent: the potential is created by the

participants' positioning and co-evolves with the price it steers. This is the precise sense in which the present model is not a parameterised extension of Black–Scholes but a change of the governing equation — from a one-way linear map to a two-way self-consistent field.

Within the stochastic-mechanical reformulation of finance (Andreev & Howden, 2026a), the spot is a configuration-space coordinate guided by a velocity field decomposed into current and osmotic parts. The dealer-flow drift enters that decomposition as a component of the current velocity sourced by the options book rather than by the price's own history; the osmotic component carries the residual order flow. Avellaneda & Lipkin (2003) is the linear, single-expiry, pinning-only limit of the present construction; the contribution here is the nonlinear, full-surface, charm-aware generalisation, together with a falsifiable forward test.

5. Pre-Registered Forward Validation

Consistent with Andreev & Howden (2026e), the model is validated on forward, point-in-time option surfaces — never on reconstructed historical signals — and the test is frozen before any outcome is observed (Harvey, Liu & Zhu, 2016). The decisive statistic is the incremental significance of the self-focusing term, not the significance of the field as a whole.

Pre-registered protocol.

(P1) Universe: SPY (primary), QQQ, IWM. Field reconstructed daily (and intraday on a live REST chain) from our own point-in-time surfaces with genuine open interest.

(P2) Regime gate: the pinning claim is tested only in the long-gamma ($\Gamma_S > 0$) regime; the short-gamma (amplifying) regime is a separate, secondary claim.

(P3) Nested regression of the signed short-horizon underlying move on the linear and incremental terms: $y = a + b \cdot v_{lin} + c \cdot (v_{nl} - v_{lin}) + \varepsilon$, Newey–West standard errors.

(P4) Decisive bar: the increment coefficient $c > 0$ with Newey–West $t > 2.0$; the mechanism bar $b > 0$ with $t > 3.0$; a deflated Sharpe ratio ≥ 0.95 charging the full multiple-testing cost (the frozen 24-combination neighbourhood: dealer sign convention $\times$ pin definition $\times$ sharpening exponent $\times$ horizon); and the predicted drift must beat a size-matched random-pin placebo (2,000 draws) at the 95th percentile.

(P5) Disqualifier: if flipping the dealer sign convention rescues the result, the sign was fit post-hoc and the result is rejected.

5.1 Base reconciliation and the frozen registration

The registration was frozen on 16 June 2026, before any outcome was observed. Beyond (P1)–(P5) it carries three base-reconciliation gates that must pass before any outcome metric is read: (A) the field reconstruction's net gamma must match an independently computed gamma-exposure capture on the same chain snapshot to within 1%; (B) the long-gamma regime must show a smaller pooled mean absolute next-day move than the short-gamma regime (the canonical stylised fact the field must reproduce);

and (C) the linear drift must be correctly signed in the long-gamma regime — if the baseline mechanism is sign-wrong, the increment test is moot. One pre-outcome amendment (of at most two permitted) is on record: the initial reconstruction merged each strike's call and put onto a single implied volatility, inflating the call-gamma leg by roughly 8% — and because net dealer gamma is a small difference of two large legs, failing gate (A) at 32.9%. Pricing each leg with its own implied volatility, a reconstruction-fidelity correction touching no frozen shape constant, restored agreement to 0.0%. The first read occurs at sixty accrued trading days per ticker and the robust read at one hundred and twenty. An offline historical accelerator was evaluated and abandoned before any scoring: the candidate free historical chain archives carry no per-strike open interest — gamma exposure is uncomputable without it — and a paid archive is out of scope for the first version, so the forward log is the sole pass-eligible path.

5.2 The accruing forward log (descriptive)

Figures 4 and 5 show the field operating on the first weeks of the forward log. They are descriptive exhibits of the reconstructed object — no regression, t-statistic, or deflated Sharpe ratio is computed, and none will be until the pre-registered quorum accrues. Their role is to document that the data feed satisfies the feasibility properties the protocol assumes: the reconstruction is non-degenerate and time-varying, both gamma regimes occur in practice, and the self-focusing modulation is a real, non-trivial transformation on live surfaces rather than a factor pinned at its clamps.




Figure 4. Reconstructed net dealer dollar gamma over the first weeks of the point-in-time forward log (12–29 June 2026), in billions of dollars per 1% spot move, for SPY, QQQ and IWM. The dashed line is the gamma flip: dealers are net long gamma above it (stabilising, pinning regime) and net short gamma below it (destabilising, amplifying regime). The reconstruction is non-degenerate and time-varying — the SPY aggregate moves from roughly $\pm\$0.1$ B early in the window to $-\$4.1$ B on 26 June and back through the flip to long gamma on 29 June — which is the feasibility property the pre-registration requires of the data feed. Descriptive only: no outcome statistic is computed. Source: the daily point-in-time surface log (per-strike open interest and per-leg implied volatilities captured after the close). Two early snapshots with incomplete chains (15 June, SPY and QQQ) are omitted.




Figure 5. The self-interaction layer on real surfaces: nonlinear drift v_{nl} against linear drift v_{lin} (both dimensionless fractional drift, $\times 10^{-3}$) for the 24 full field records accrued 18–29 June 2026. The dashed identity line is the no-self-modulation locus ($\Phi = 1$): points below it are records where the spot sat away from the concentrated-gamma strikes and the local-density factor attenuated the linear pull (down to its frozen clamp floor of 0.20); points above it are self-trapping records where the local factor exceeded one (up to 2.86). In every accrued record the expiry-sharpening factor sits at its clamp floor of 1 (the tracked tenor is near the one-month reference), so the modulation shown is entirely the local-density channel. The vertical displacement from the identity line is exactly the self-focusing increment $v_{nl} - v_{lin}$ whose incremental predictive power is the decisive pre-registered bar (P4); no such statistic is computed here.

Caveats. The end-of-day horizon is the weakest test of a mechanism that operates intraday; a daily null does not refute the intraday hypothesis but redirects it. That redirection should be stated as a property of this study rather than as a contingency, because the regime invoked in Section 1 as the model's motivation is the zero-days-to-expiry dealer-gamma regime, and the forward log accruing under this protocol is an end-of-day daily capture which does not observe it: the pinning, charm and vanna channels that make the regime interesting resolve inside the session the log samples once. The registered forward test is therefore a test of the field's daily-horizon residue, and a null on it redirects to a separately registered intraday hypothesis rather than counting as evidence against the mechanism. The registration forbids the alternative response of re-rolling the end-of-day threshold until it fires. The dealer sign convention is an assumption tested as a disqualifier, not a free parameter. The self-focusing clamps are frozen safeties whose sensitivity

is reported. The field is a microstructure object and is excluded by construction from any portfolio-allocation decision.

6. Discussion

The contribution is a change of governing equation, not a parameter. Black–Scholes assumes the underlying is exogenous; the empirical record in the dealer-gamma regime says it is not. By identifying the dealer-gamma landscape with a quantum potential, the hedging flow with a guidance drift, and the configuration-dependence of the local gamma density with a Gross–Pitaevskii self-interaction, we obtain a single field model whose linear limit is the practitioner GEX statistic and whose nonlinear term is a falsifiable, forward-tested object. A confirmation would be a genuinely new pricing/mechanics result — a model of the spot–options feedback that classical option theory is structurally unable to represent. A null at the daily horizon would localise the search to

intraday resolution. Either outcome is informative, and neither is asserted here: this is a methods and protocol paper, with empirical results to be reported separately and only on forward data.

References

- Almgren, R. & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.
- Avellaneda, M. & Lipkin, M. D. (2003). A market-induced mechanism for stock pinning. *Quantitative Finance*, 3(6), 417–425.
- Baaquie, B. E. (2004). *Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates*. Cambridge University Press.
- Barbon, A. & Buraschi, A. (2020). Gamma fragility. *Swiss Finance Institute Research Paper*, 20–25.
- Black, F. & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Bohm, D. (1952). A suggested interpretation of the quantum theory in terms of “hidden” variables, I. *Physical Review*, 85(2), 166–179.
- Gârleanu, N., Pedersen, L. H. & Potesman, A. M. (2009). Demand-based option pricing. *Review of Financial Studies*, 22(10), 4259–4299.
- Gross, E. P. (1961). Structure of a quantized vortex in boson systems. *Il Nuovo Cimento*, 20(3), 454–477.
- Harvey, C. R., Liu, Y. & Zhu, H. (2016). ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
- Madelung, E. (1927). Quantentheorie in hydrodynamischer Form. *Zeitschrift für Physik*, 40(3–4), 322–326.
- Nelson, E. (1966). Derivation of the Schrödinger equation from Newtonian mechanics. *Physical Review*, 150(4), 1079–1085.
- Ni, S. X., Pearson, N. D. & Potesman, A. M. (2005). Stock price clustering on option expiration dates. *Journal of Financial Economics*, 78(1), 49–87.
- Pitaevskii, L. P. (1961). Vortex lines in an imperfect Bose gas. *Soviet Physics JETP*, 13(2), 451–454.
- Andreev, M. A. & Howden, J. (2026a). Stochastic mechanical methods for quantitative portfolio management. *Quark Quantitative Research Working Paper*.
- Andreev, M. A. & Howden, J. (2026e). Ontological limits of historical validation for wavefunction-calibrated financial systems. *Quark Quantitative Research Working Paper*.